arXiv:1602.02950v1 [cs.LG] 9 Feb 2016

Spoofing detection under noisy conditions: a preliminary investigation and an initial database

Xiaohai Tian[1,2], *Zhizheng Wu*[3], *Xiong Xiao*[4], *Eng Siong Chng*[1,2,4] *and Haizhou Li*[1,5]

[1]School of Computer Engineering, Nanyang Technological University (NTU),Singapore
[2]Joint NTU-UBC Research Center of Excellence in Active Living for the Elderly, NTU, Singapore
[3]The Center for Speech Technology Research, University of Edinburgh, United Kingdom
[4]Temasek Laboratories, NTU, Singapore
[5]Human Language Technology Department, Institute for Infocomm Research, Singapore
{xhtian, xiaoxiong, aseschng}@ntu.edu.sg, zhizheng.wu@ed.ac.uk, hli@i2r.a-star.edu.sg

Abstract

Spoofing detection for automatic speaker verification (ASV), which is to discriminate between live speech and attacks, has received increasing attentions recently. However, all the previous studies have been done on the clean data without significant additive noise. To simulate the real-life scenarios, we perform a preliminary investigation of spoofing detection under additive noisy conditions, and also describe an initial database for this task. The noisy database is based on the ASVspoof challenge 2015 database and generated by artificially adding background noises at different signal-to-noise ratios (SNRs). Five different additive noises are included. Our preliminary results show that using the model trained from clean data, the system performance degrades significantly in noisy conditions. Phase-based feature is more noise robust than magnitude-based features. And the systems perform significantly differ under different noise scenarios.

1. Introduction

Recently, automatic speaker verification (ASV) has been significantly advanced to the point of mass-market adoption [1–4]. However, most of current ASV systems assume human voices, and there are concerns that whether the systems can still achieve robust performance in the face of diverse spoofing attacks. A spoofing attack is that an attacker attempts to manipulate an ASV result for a target genuine speaker to obtain access permission. A significant amount of evidences have confirmed the vulnerability of current state-of-the-art ASV systems under spoofing attacks as reviewed in [5]. This has led to the active development of spoofing countermeasures, also called spoofing detection, that is to discriminate human and spoofed speech.

In the past several years, spoofing detection for speaker recognition has been studied on a variety of diverse datasets. In [6, 7], the Wall Street Journal (WSJ) corpus was used to assess countermeasures for speech synthesis attacks. In [8], the publicly available RSR2015 corpus was used to evaluate spoofing detection for replay attacks. In [9, 10], synthetic speech from the Blizzard challenge [11] was used for speech synthesis spoofing detection. In [12], a recently released spoofing and anti-spoofing (SAS) corpus as a standard spoofing database was used to assess speech synthesis and voice conversion spoofing countermeasures. We note that WSJ, SAS and Blizzard challenge databases were recorded by high-quality microphones in sound-proofing environment, while the RSR2015 corpus was recorded by multiple mobile devices in a quiet office room. All these databases do not have any significant channel and/or additive noise. These databases allow us to focus on spoofing effects but do not simulate practical scenarios of ASV applications.

There are also some studies that use data with channel noise. The National Institute of Standards and Technology (NIST) Speaker Recognition Evaluation (SRE) 2006 database which has significant telephone channel noise was used to assess voice conversion spoofing countermeasures in [13–17]. In [18], a so-called AVspoof database includes replay, speech synthesis and voice conversion spoofing attacks to simulate realistic scenarios, which re-recorded synthetic or voice-converted speech using multiple mobile devices.

In general, all the databases used in the past spoofing detection studies do not consider additive noise, even the recent standard spoofing detection databases: SAS[1] and ASVspoof 2015 challenge[2] databases. However, in practical scenarios, it is hard to avoid additive and/or channel noise. Hence, another concern for ASV deployment arises that whether currently developed spoofing detection algorithms/systems are still effective under noisy conditions.

In this work, we focus on spoofing detection under additive noisy conditions. We perform a preliminary investigation of spoofing detection under noisy conditions using current state-of-the-art countermeasure techniques; and also briefly introduce an initial database we built for the task[3]. In general, we aim to answer the following questions:

- Do current state-of-the-art spoofing detection algorithms work well under additive noisy conditions?
- How additive noises affect the spoofing detection performance?
- What kind of noise is more serious to degrade the performance of spoofing detection algorithms?

We believe better understanding of above questions and the noisy database will drive the development of generalised and noise robust spoofing detection algorithms.

[1]SAS corpus is available at: http://dx.doi.org/10.7488/ds/252

[2]ASVspoof 2015 corpus is available at: http://dx.doi.org/10.7488/ds/298

[3]The noisy database will be publicly-available for free under a CC-BY license.

2. Noisy Database

In order to represent the practical application scenarios for spoofing detection, we attempt to design a database in additive noisy environments based on the ASVspoof 2015 challenge database [19].

ASVspoof database is a spoofing and anti-spoofing database, consisting of both genuine (human speech) and ten types of spoofed speech (named as S1-S10 in ASVspoof 2015 challenge) implemented by three speech synthesis and seven voice conversion spoofing algorithms. The ASVspoof database contains three subsets, including training, development and evaluation sets. The training and development sets only contain *known* attacks (S1-S5); while the evaluation set consists of both *known* and *unknown* (S6-S10) attacks. More details and protocols about the ASVspoof database can be found in [19].

This noisy version aims to quantify the effects of current spoofing detection algorithms in noise conditions and to facilitate future assessment work in this task. In this section, we will briefly introduce the types of noise to be added, and the procedure of adding noise.

2.1. Noise signals

Five types of noise signals, representing the probable application scenarios, are used for the construction of the noisy ASVspoof database. A subset of three types of noise, white noise, speech babble and vehicle interior noise, are selected from NOISEX-92 database [20]. Another two types mixed noise, street noise and cafe noise, are selected from QUT-NOISE database [21]. These are standard types of additive noise used for speech recognition [21–23], speaker verification [24, 25] and speech enhancement [26]. We briefly describe these noises as follows:

- **White Noise**: The random signal with a constant power spectral density.
- **Babble Noise**: Speech babble and the recording is made in a canteen with 100 people speaking.
- **Volvo Noise**: Vehicle interior noise and the recording is made in Volvo 340 on an asphalt road, rainy conditions.
- **Street Noise**: Mixed noise, which is made at the roadside near inner-city, mainly consisting of road traffic noise, pedestrian traffic noise and bird noise.
- **Cafe Noise**: Mixed noise, which is made in outdoor cafe environment, mainly consisting of speech babble and kitchen noise from the cafe environment.

Figure 1 shows the spectrogram of all the five noises. We can classify the noises into stationary noise and non-stationary noise. White noise and Volvo noise are stationary noise. While babble noise, street noise and cafe noise are recorded in non-stationary noise environment, whose the magnitude and phase spectrogram are varied over time.

2.2. Adding noise

The data from ASVspoof database are considered as clean data. Noise is artificially added to the clean data. The Filtering and Noise Adding Tool (FaNT)[4] is used for the adding noise process. The noisy signals are generated by adding the clean speech and noise files together at various SNRs. As the silence periods appear in many speech files of ASVspoof database, it

[4]http://dnt.kr.hs-niederrhein.de/

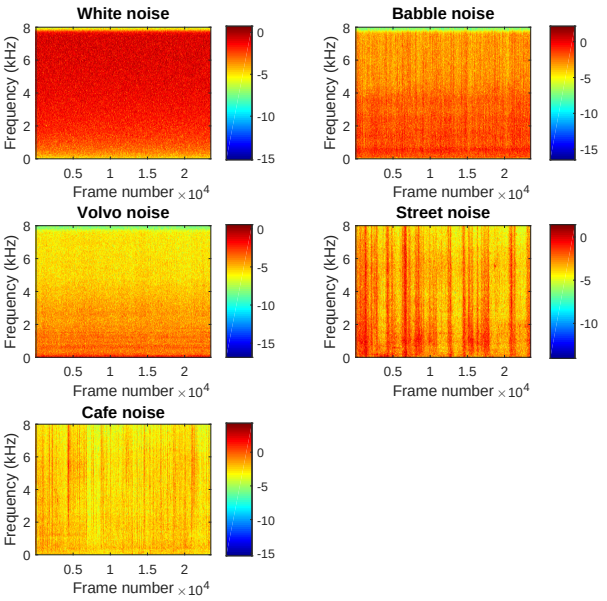

Figure 1: The spectrogram of different noise file.

is important to calculate the SNR only based on the sections of speech signal. The bandpass filter or frequency weighting are often used for SNR calculation, to ensure the SNR are appropriate and comparable. In this work, to add the noise based on the human hearing perception, we define the SNR as the ratio of signal to noise energy after filtering both signals with the A-weighting filter. This filter emphasizes the frequencies around 3 kHz to 6 kHz where the human ear is most sensitive and give a lower response for the very high and very low frequencies to which the ear is less sensitive [27]. Hence, noises such as Volvo noise, whose energy distribution concentrates more on very low frequency (below 1 kHz), tend to have higher energy at the same SNR level.

For each clean signal in the development and evaluation sets of ASVspoof database, fifteen noisy versions of the signal are generated consisting of five types of noise in three SNR levels. Given a clean signal, for each noise type, we take a segment of the noise signal with equal length as the clean signal but random starting point from the whole noise file. Then the noise segment is scaled and added to the clean signal in 20 dB, 10 dB and 0 dB SNR levels.

After the adding noise process, the clipping may occur, especially at low SNR levels due to high noise energy. In order to maintain a stable spectrogram representation of the signal, the signal is scaled to avoid the clipping.

All the five types of noise are added to the ASVspoof database at the SNRs of 20 dB, 10dB, and 0 dB. Hence, the ASVspoof noisy database is fifteen times of the clean database.

3. Benchmarking system

In order to demonstrate the utility of the ASVspoof noisy database, we conduct a series of experiments to examine the performance of spoofing speech detection system on a range of SNRs in all five noise scenarios.

The detection system, as shown in Figure 2, consists of (a) the feature extraction module to extract six types of feature used

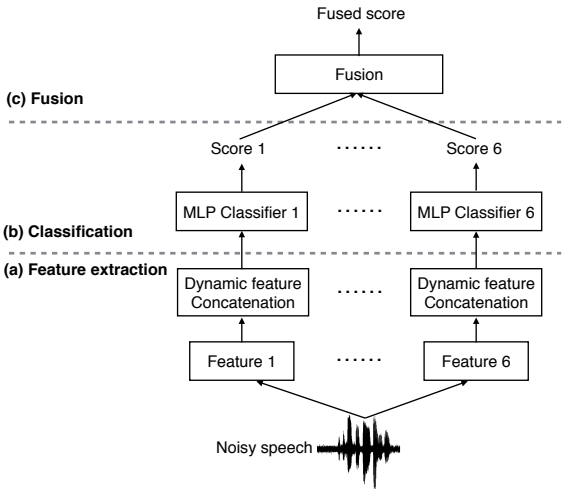

Figure 2: The architecture of our detection system, including (a) the feature extraction module, (b) the classification module and (c) the score fusion module.

for classification; (b) the classification module to calculate the score for each feature; (c) the score fusion module to fused the scores obtained from six classifiers. The details of these three modules are introduced as follows.

3.1. Feature extraction

Similar to our previous system described in [28, 29], six types of feature are extracted. As shown in Figure 2 (a), given a noisy waveform, the Hamming window and direct current (DC) offset removal are applied on each analysis frame. Then short-time Fourier transform (STFT) is applied on the speech signal using analysis window of 25ms with 15ms overlap. The FFT length is chosen to be 512 and the dimension of all the original features are 256. For the n-th frame, the magnitude and phase spectrum, $|\mathrm{X}(n,\omega)|$ and $\theta(n,\omega)$, are obtained by

$$\mathrm{X}(n,\omega) = |\mathrm{X}(n,\omega)|e^{j\theta(n,\omega)}, \tag{1}$$

After that, two magnitude-based features, namely log magnitude spectrum (LMS) and residual log magnitude spectrum (RLMS) are derived from magnitude spectrum. Four phase-based features, namely instantaneous frequency derivative (IF), baseband phase difference (BPD), group delay (GD) and modified group delay (MGD), are derived from phase spectrum.

The features are summarized as follows:

- **LMS**: The log magnitude spectrum feature can be expressed as $\mathrm{LMS(n},\omega) = log(|\mathrm{X}(n,\omega)|)$. The LMS contains the formant information, harmonic structure and all the spectral detail of speech signal.
- **RLMS**: The residual log magnitude spectrum feature is the LMS extracted from the linear predictive coding (LPC) residual signal. As the formant information is removed, this feature can better analyses the harmonic structure and spectral details.
- **IF**: Instantaneous frequency [30] is the derivative of the phase along time axis and captures the temporal information of phase. It is defined as:

$$\mathrm{IF}(n,\omega) = \mathrm{princ}(\theta(n,\omega) - \theta(n-1,\omega)). \tag{2}$$

where $\mathrm{princ}(\cdot)$ represents the principal value operator, mapping the input onto $[-\pi;\pi]$ interval by adding integer numbers of 2π.

- **BPD**: Baseband phase difference [31] is another phase feature derived from IF and baseband STFT. For the n-th frame, the BPD is calculated as:

$$\mathrm{BPD}(n,\omega) = \mathrm{princ}(\mathrm{IF}(n,\omega) - \Omega_k l), \tag{3}$$

where $\Omega_k = 2\pi k/L$ is the normalized angular frequencies.

- **GD**: Group delay [32] is a representation of filter phase response, which is defined as the negative derivative of the Fourier transform phase. It is a frame-based feature, used to capture the phase distortion along frequency axis.

$$\mathrm{GD}(n,\omega) = \mathrm{princ}\{\theta(n,\omega) - \theta(n,\omega-1)\}, \tag{4}$$

- **MGD**: A variation of GD, which can obtain a more clear phase pattern than GD. The MGD [32] feature of frame n is calculated as:

$$\tau(n,\omega) = \frac{X_R(n,\omega)Y_R(n,\omega) + X_I(n,\omega)Y_I(n,\omega)}{|S(n,\omega)|^{2\gamma}}, \tag{5}$$

$$\mathrm{MGD}(n,\omega) = \frac{\tau(n,\omega)}{|\tau(n,\omega)|}|\tau(n,\omega)|^{\alpha}, \tag{6}$$

where $X_R(n,\omega)$ and $X_I(n,\omega)$ denote the real and image part of STFT for $x(l)$; while $Y_R(n,\omega)$ and $Y_I(n,\omega)$ denote the real and image part of STFT for $lx(l)$, respectively. $S(n,\omega)$ is the smoothed spectrum of $|X(n,\omega)|$. Based on experimental results, the values of γ and α are set to 0.7 and 0.2 respectively.

In Fig 3 and Fig 4, we present the LMS and IF features extracted from clean and the noisy signals. The noisy signals include the white and street noise scenarios at the SNR of 20 dB, 10 dB and 0 dB. We observe that both LMS and IF feature are distorted by additive noise significantly. The patterns become more blurred for lower SNRs.

3.2. Classifier

Figure 2 (b) shows the classification part of the detection system. Our previous multilayer perceptron (MLP) based spoofing speech detection system [29] is used in this work. Each of the features mentioned above with its delta and acceleration coefficients is used as the input vector to train its own classifier. The MLP, which contains one hidden layer with 2,048 sigmoid nodes, is used to predict the posterior probability of the input vector being synthetic speech. The score is calculated by averaging the posterior probabilities of all the frames over the utterance. Noted that, in this work, all the MLP classifier are trained from clean data.

3.3. Evaluation metrics and fusion

The equal error rate (EER), where false acceptance rate and miss rejection rate become equal, is used to evaluate the system performance.

As described in Section 3.1, different features are designed to detect different types of artifacts. In order to benefit the advantages of each feature and improve the system stability, a score level fusion is applied. Figure 2 (c) shows score fusion of the detection system.

The fusion is applied on the feature-based results in different noise scenarios. Based on our preliminary experiments, the

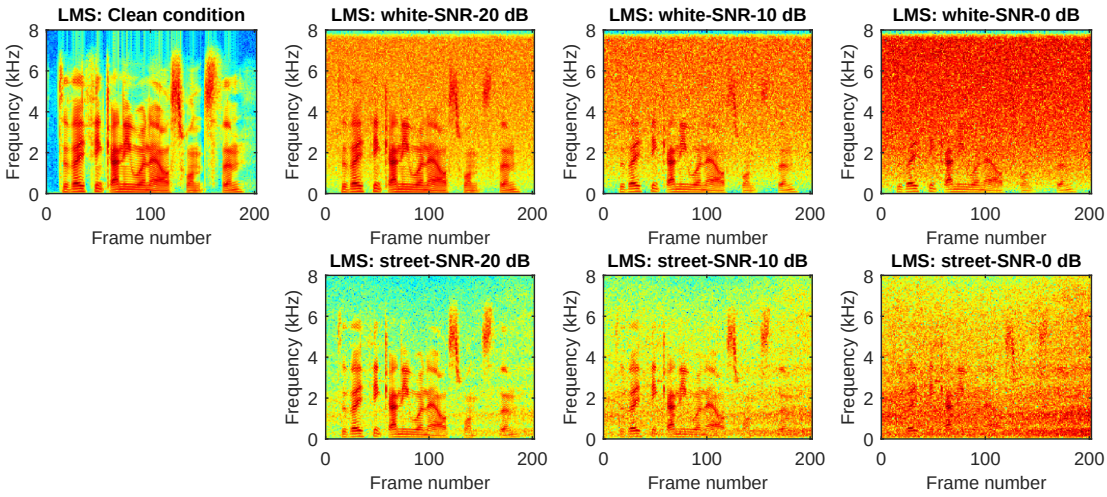

Figure 3: Demonstration of the LMS feature for utterance D14_1000302, in both clean and noise scenarios.

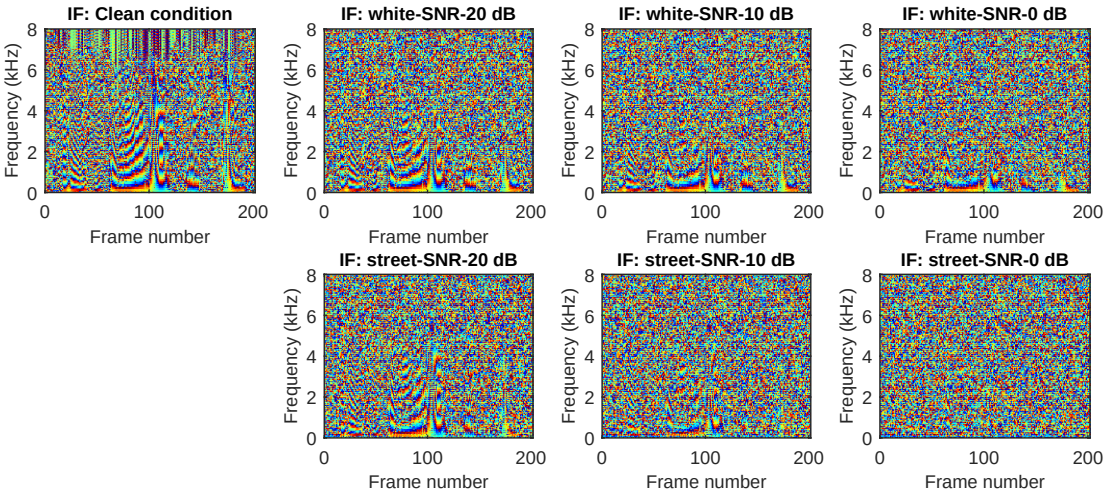

Figure 4: Demonstration of the IF features for utterance D14_1000302, in both clean and noise scenarios.

weighted summation fusion, tuned on the development data set, exhibits the over-fitting effects at street and cafe noise scenarios. To avoid this problem, the scores of all systems are simply averaged to produce the final score.

In this work, the Bosaris toolkit[5] is used to compute the EERs of each feature and the fused system.

4. Experiments

4.1. Experimental setups

The database used in the experiments consist of three subsets, including training set, development set and evaluation set. The training set is clean speech data taken from ASVspoof database. As the training set consists of clean data only, it models the speech without noise distortion and represents all the speech information. The best performance of the clean classifier is obtained in the case of testing on clean data, which can be found in our previous work [29].

The development and evaluation sets are chosen from the noisy ASVspoof database, including five noise scenarios at three different SNRs as described in Section 2. Because the classifier used in these experiments is the same as that of our previous work [29], these results are comparable with the results in clean condition.

[5]https://sites.google.com/site/bosaristoolkit/

4.2. Evaluation results

As the results on the development set are similar to that of the S1 to S5 on the evaluation set, only the feature-based results of the evaluation set are reported. The results are shown independently as the known attacks (S1-S5), the unknown attacks (S6-S9) and the unknown attacks generated by waveform concatenation (S10). The EERs are listed in Table 1 for both noisy dataset and clean dataset using the clean classifier.

We first analyse the effect of noisy data for the detection system using different features. In general, across all the five noise scenarios, the systems perform worse than that of clean condition. As expected, in most spoof attacks, the detection performance deteriorates as SNR decreases. We notice that, in most noisy scenarios, the magnitude-based features, LMS

Table 1: *Average EERs (%) of different features on the evaluation set. Clean indicates the results of our previous work [29].*

	S1-S5 (Known)							S6-S9 (Unknown)							S10 (Unknown)						
Feature	LMS	RLMS	IF	BPD	GD	MGD	Fusion	LMS	RLMS	IF	BPD	GD	MGD	Fusion	LMS	RLMS	IF	BPD	GD	MGD	Fusion
clean	**0.02**	**0.34**	**1.31**	**0.10**	**0.05**	**0.00**	**0.00**	**0.01**	**0.36**	**1.31**	**0.09**	**0.03**	**0.02**	**0.00**	**35.24**	**30.80**	**25.56**	**30.67**	**33.90**	**38.54**	**28.87**
white_snr_20	2.49	33.94	3.35	3.24	6.52	8.70	2.50	7.45	35.35	4.23	3.64	5.86	11.65	3.20	46.02	48.42	39.23	37.68	46.40	47.72	43.36
white_snr_10	23.77	34.46	5.89	4.70	12.04	24.55	3.74	15.40	35.71	8.42	6.34	11.60	26.61	4.86	48.07	48.81	41.47	39.70	47.62	47.72	45.22
white_snr_0	37.80	37.85	14.01	11.66	24.49	42.89	10.64	32.17	41.09	20.00	17.16	25.69	43.18	15.27	48.31	49.16	45.00	44.32	48.85	47.71	46.80
babble_snr_20	7.64	47.38	3.45	5.61	4.82	2.93	3.77	7.29	48.04	4.13	6.66	5.72	4.64	4.71	45.98	48.41	40.12	42.35	44.34	47.68	46.11
babble_snr_10	15.40	48.50	12.01	14.81	17.28	8.58	10.09	14.99	48.49	15.61	19.49	20.09	13.94	14.11	42.06	48.22	45.18	45.60	44.96	47.71	46.44
babble_snr_0	28.19	48.42	29.45	31.80	34.48	25.06	25.09	25.72	48.52	33.78	36.78	36.29	31.54	28.53	43.46	48.34	46.50	46.84	46.45	47.62	47.19
volvo_snr_20	13.75	42.53	15.48	19.69	16.55	15.34	9.69	14.59	47.39	19.10	21.76	19.86	19.71	12.12	47.79	44.70	42.73	47.75	47.26	47.89	46.16
volvo_snr_10	27.22	45.21	26.90	29.04	30.37	19.57	20.61	27.63	47.21	30.76	33.06	34.08	24.15	23.93	48.08	47.28	44.52	47.85	47.41	47.00	47.88
volvo_snr_0	38.89	47.56	37.28	38.05	40.84	30.37	37.98	39.69	47.07	39.56	41.00	43.21	34.54	39.24	47.88	49.12	45.84	47.68	47.30	46.54	47.94
street_snr_20	15.92	39.96	6.49	7.55	11.61	9.61	7.87	13.81	38.73	7.80	9.82	15.43	13.90	8.38	37.70	44.22	42.59	43.79	45.38	47.62	37.97
street_snr_10	30.17	44.60	13.72	13.95	22.91	20.52	18.65	26.15	44.82	15.77	18.00	26.51	25.89	18.72	41.02	47.69	45.25	45.75	45.38	47.78	39.65
street_snr_0	43.44	47.35	27.27	28.52	34.80	33.44	32.44	39.22	47.03	28.78	32.33	35.81	37.17	31.65	48.80	50.00	47.05	46.94	47.18	47.88	47.53
cafe_snr_20	12.32	34.76	4.00	5.19	8.21	7.23	4.64	8.70	31.25	4.39	6.50	9.16	10.14	4.15	38.54	48.32	40.82	40.18	46.14	47.77	41.05
cafe_snr_10	26.53	42.61	5.91	6.83	15.91	20.29	16.16	20.07	40.09	6.86	9.09	15.88	23.88	13.91	42.10	50.00	41.70	42.34	45.88	47.77	44.35
cafe_snr_0	44.79	46.97	13.72	13.34	28.63	39.28	35.39	39.46	45.55	15.00	17.32	25.67	39.43	32.16	49.92	50.00	45.21	44.96	47.33	47.82	48.51
Average EER cross noise scenarios	24.55	42.81	14.60	15.60	20.63	20.56	15.95	22.16	43.09	16.95	18.60	22.06	24.02	17.00	45.05	48.18	43.55	44.25	46.53	47.62	45.08

and RLMS, perform worse than the phase-based features, IF, BPD, GD and MGD. In particular, in all the spoofing attacks, the RLMS obtains much higher EERs than other features. This may be due to that the LPC filter is not robust in noisy environments [33], which affects the quality of RLMS. Among the phase-based features, IF and BPD outperform other features in terms of the average EERs over all the noise scenarios. We also found that, some features are effective for particular noise scenario. For example, in the babble noise scenario, the MGD is capable to obtain low error rates. While in white, street, and cafe noises, IF and BPD perform much better than other features.

Next, we compare the performance across different types of attack in noise conditions. Because S1-S5 attacks are available for training, even in noisy condition, the lower error rates are obtained in these attacks than the rest types of attack. Although the error rates of S6-S9 is higher than that of S1-S5, the results still comparable. This is consistent with the results in clean condition [29]. For S10, even apply the score fusion, the error rates of all the features are significantly higher than that of S1-S9. Hence, we conclude that in both clean and noisy conditions, the detection of S10 is still the most challenge task among the spoofing attacks.

4.3. Fusion results

Table 2 presents the results of fused systems on both development and evaluation sets. We first examine the results at different SNRs. In all the noise scenarios at the SNR of 20 dB, the system performance degradation significant differ for different noise scenarios. On the development set, the EERs of different fuse systems are between 2.52% (cafe noise) to 8.32% (Volvo noise); while on the evaluation set, the EERs are varies between 5.25% (white noise) to 14.31% (Volvo noise). However, at the SNR of 0 dB, all the systems performance degrade significantly on both development and evaluation set. Fig 3 and Fig 4 show that most of the feature patterns are lost in such low SNR.

Then, we analyse the fused results in different noise scenarios. Among all the noise scenarios, the system under white noise scenario performs best, which constantly achieves lowest error rates. Especially, at low SNRs, 10 dB and 0 dB, the system under white noise outperform that of other noise scenarios significantly. As discussed in Section 2.2, compare to other types of noise, the Volvo noise tends to give higher energy at the same SNR level, resulting in more distortions in the noisy signals and higher EERs.

For the non-stationary noisy conditions, the features distorted by such noise are time-varying. Consequently, in these noise conditions, the system performance degrades more than white noise scenarios. This provides more challenge to the system to detect the spoofed attacks in such noisy conditions.

5. Conclusions

In this paper, we constructed a noisy database for spoofing and anti-spoofing research. This database is generated from the ASVspoof database, adding five types of additive noise at three SNR levels. To provide the benchmark results, we use the state-of-the-art spoofing detection system to detect the spoofing attacks in noisy conditions. The preliminary results using the classifier trained from clean data shown that,

- the performance of detection systems degrade in all the noise scenarios. The system performance deteriorates as SNR decreases;
- in noisy environments, even the best EER is about 4% lower (white noise at 20 dB) than that of clean condition;
- in general, the non-stationary noises affect the detection system more seriously;
- the system performance varies significantly under different noise scenarios and the phase-based features are noise robust than magnitude-based features.

In this paper, we only presented benchmark results to demonstrate the vulnerability of current spoofing detection systems under additive noise conditions. In future work, we will use convolutional noise, such as channel noise and reverberate noise, to simulate the spoofing attacks in more complex noise scenarios. Moreover, the classifier presented in this work was trained from the clean data. There are also plans to exam the effectiveness of multi-condition training for spoofing detection under noisy conditions. Finally, a proper score fusion method will be investigated to both avoid the over-fitting effects and improve the system performance.

Table 2: *EERs (%) of fused system on both development and evaluation sets. Clean indicates the results of our previous work [29].*

	Development						Evaluation										
Feature	S1	S2	S3	S4	S5	Average	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	Average
clean	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**28.87**	**2.89**
white_snr_20	2.49	4.23	0.56	0.51	6.47	2.85	2.77	3.31	0.32	0.31	5.81	5.73	2.86	1.81	2.40	43.36	6.87
white_snr_10	3.92	6.33	2.42	2.21	7.40	4.46	4.38	5.18	1.39	1.32	6.41	6.40	5.61	2.98	4.46	45.22	8.34
white_snr_0	10.30	16.04	6.27	5.93	14.55	10.62	11.99	16.00	5.31	5.09	14.83	15.67	18.29	12.66	14.45	46.80	16.11
babble_snr_20	7.17	5.97	2.42	2.38	8.42	5.28	4.65	4.48	0.77	0.73	8.25	8.75	3.60	3.65	2.84	46.11	8.38
babble_snr_10	16.51	17.14	5.02	4.70	14.30	11.53	14.65	15.63	2.76	2.51	14.90	17.49	18.14	10.44	10.40	46.44	15.33
babble_snr_0	32.48	33.45	15.88	15.83	29.96	25.52	31.75	32.93	15.04	14.90	30.84	33.34	35.63	18.76	26.40	47.19	28.68
volvo_snr_20	14.21	12.94	0.86	0.77	12.80	8.32	19.66	14.46	1.46	1.47	11.37	12.35	15.03	14.12	6.97	46.16	14.31
volvo_snr_10	31.83	27.06	10.80	10.60	27.86	21.63	33.39	24.96	9.68	9.55	25.46	24.86	27.16	26.41	17.30	47.88	24.67
volvo_snr_0	44.42	40.40	31.53	31.59	43.06	38.20	45.14	39.85	31.87	31.91	41.15	39.98	39.74	40.69	36.55	47.94	39.48
street_snr_20	4.87	6.42	1.97	1.77	7.44	4.49	8.72	10.75	5.20	5.28	9.38	11.03	10.52	4.90	7.07	37.97	11.08
street_snr_10	13.23	15.99	8.46	8.16	14.32	12.03	19.94	22.64	15.26	15.27	20.13	22.52	22.34	12.02	18.01	39.65	20.78
street_snr_0	26.62	29.18	23.01	23.12	27.75	25.94	33.20	34.53	30.52	30.48	33.47	34.42	34.26	26.96	30.95	47.53	33.63
cafe_snr_20	1.56	3.00	1.00	0.83	6.21	2.52	4.28	4.77	3.69	3.62	6.84	7.14	4.05	2.03	3.37	41.05	8.08
cafe_snr_10	6.12	8.69	6.25	6.13	9.67	7.37	14.15	17.06	16.76	16.74	16.11	18.28	16.53	7.68	13.14	44.35	18.08
cafe_snr_0	25.26	26.73	29.01	28.66	27.37	27.41	33.33	33.51	38.41	38.48	33.23	33.92	33.44	29.73	31.55	48.51	35.41

6. Acknowledgment

This research is supported by the National Research Foundation Singapore under its Interactive Digital Media (IDM) Strategic Research Programme. This work is also supported by the DSO funded project MAISON DSOCL14045, Singapore.

7. References

[1] Kong Aik Lee, Bin Ma, and Haizhou Li, "Speaker verification makes its debut in smartphone," in *IEEE Signal Processing Society Speech and language Technical Committee Newsletter*, February 2013.

[2] Mikhail Khitrov, "Talking passwords: voice biometrics for data access and security," *Biometric Technology Today*, vol. 2013, no. 2, pp. 9–11, 2013.

[3] Brett Beranek, "Voice biometrics: success stories, success factors and what's next," *Biometric Technology Today*, vol. 2013, no. 7, pp. 9–11, 2013.

[4] Wenchao Meng, Duncan Wong, Steven Furnell, and Jianying Zhou, "Surveying the development of biometric user authentication on mobile phones," *IEEE Communications Surveys and Tutorials*, 2015.

[5] Zhizheng Wu, Nicholas Evans, Tomi Kinnunen, Junichi Yamagishi, Federico Alegre, and Haizhou Li, "Spoofing and countermeasures for speaker verification: a survey," *Speech Communication*, vol. 66, pp. 130–153, 2015.

[6] P. L. De Leon, M. Pucher, J. Yamagishi, I. Hernaez, and I. Saratxaga, "Evaluation of speaker verification security and detection of HMM-based synthetic speech," *IEEE Trans. Audio, Speech and Language Processing*, vol. 20, no. 8, pp. 2280–2290, 2012.

[7] Zhizheng Wu, Xiong Xiao, Eng Siong Chng, and Haizhou Li, "Synthetic speech detection using temporal modulation feature," in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2013.

[8] Zhizheng Wu, Sheng Gao, Eng Siong Chng, and Haizhou Li, "A study on replay attack and anti-spoofing for text-dependent speaker verification," in *Proc. Asia-Pacific Signal Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2014.

[9] Phillip L De Leon, Bryan Stewart, and Junichi Yamagishi, "Synthetic speech discrimination using pitch pattern statistics derived from image analysis," in *INTERSPEECH*, 2012.

[10] Jon Sanchez, Ibon Saratxaga, Inma Hernaez, Eva Navas, Daniel Erro, and Tuomo Raitio, "Toward a universal synthetic speech spoofing detection using phase information," *IEEE Trans. on Information Forensics and Security*, vol. 10, no. 4, pp. 810–820, 2015.

[11] Simon King, "Measuring a decade of progress in text-to-speech," *Loquens*, vol. 1, no. 1, pp. e006, 2014.

[12] Zhizheng Wu, Phillip L. De Leon, Cenk Demiroglu, Ali Khodabakhsh, Simon King, Zhen-Hua Ling, Daisuke Saito, Bryan Stewart, Tomoki Toda, Mirjam Wester, and Junichi Yamagishi, "Anti-spoofing for text-independent speaker verification: An initial database, comparison of countermeasures, and human performance," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 2016.

[13] Zhizheng Wu, Eng Siong Chng, and Haizhou Li, "Detecting converted speech and natural speech for anti-spoofing attack in speaker recognition," in *Proc. Interspeech*, 2012.

[14] Zhizheng Wu, Tomi Kinnunen, Eng Siong Chng, Haizhou Li, and Eliathamby Ambikairajah, "A study on spoofing attack in state-of-the-art speaker verification: the telephone speech case," in *Proc. Asia-Pacific Signal Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2012.

[15] Federico Alegre, Asmaa Amehraye, and Nicholas Evans, "Spoofing countermeasures to protect automatic speaker verification from voice conversion," in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2013.

[16] Elie Khoury, Tomi Kinnunen, Aleksandr Sizov, Zhizheng Wu, and Sébastien Marcel, "Introducing i-vectors for joint anti-spoofing and speaker verification," in *Proc. Interspeech*, 2014.

[17] Aleksandr Sizov, Elie Khoury, Tomi Kinnunen, Zhizheng Wu, and Sebastien Marcel, "Joint speaker verification and antispoofing in the i-vector space," *IEEE Trans. on Information Forensics and Security*, vol. 10, no. 4, pp. 821–832, 2015.

[18] Serife Kucur Ergunay, Elie Khoury, Alexandros Lazaridis, and Sebastien Marcel, "On the vulnerability of speaker

verification to realistic voice spoofing," in *IEEE International Conference on Biometrics Theory, Applications and Systems (BTAS)*, 2015.

[19] Zhizheng Wu, Tomi Kinnunen, Nicholas Evans, Junichi Yamagishi, Cemal Hanilçi, Md Sahidullah, and Aleksandr Sizov, "ASVspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Proc. Interspeech*, 2015.

[20] A. P. Varga, H. J. M. Steeneken, M. Tomlinson, and D. Jones, "The NOISEX-92 study on the effect of additive noise on automatic speech recognition," in *Tech. Rep., DRA Speech Research Unit*, 1992.

[21] David Dean, Ahilan Kanagasundaram, Houman Ghaemmaghami, Md Hafizur Rahman, and Sridha Sridharan, "The QUT-NOISE-SRE protocol for the evaluation of noisy speaker recognition," in *Proc. Interspeech*, 2015.

[22] Andrew Varga and Herman JM Steeneken, "Assessment for automatic speech recognition: II. NOISEX-92: A database and an experiment to study the effect of additive noise on speech recognition systems," *Speech communication*, vol. 12, no. 3, pp. 247–251, 1993.

[23] Hans-Günter Hirsch and David Pearce, "The aurora experimental framework for the performance evaluation of speech recognition systems under noisy conditions," in *ASR2000-Automatic Speech Recognition: Challenges for the new Millenium ISCA Tutorial and Research Workshop (ITRW)*, 2000.

[24] Andrzej Drygajlo and Mounir El-Maliki, "Speaker verification in noisy environments with combined spectral subtraction and missing feature theory," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 1998.

[25] Rahim Saeidi, Jouni Pohjalainen, Tomi Kinnunen, and Paavo Alku, "Temporally weighted linear prediction features for tackling additive noise in speaker verification," *IEEE Signal Processing Letters*, vol. 17, no. 6, pp. 599–602, 2010.

[26] Nathalie Virag, "Single channel speech enhancement based on masking properties of the human auditory system," *IEEE Transactions on Speech and Audio Processing*, vol. 7, no. 2, pp. 126–137, 1999.

[27] Harvey Fletcher and Wilden A Munson, "Loudness, its definition, measurement and calculation*," *Bell System Technical Journal*, vol. 12, no. 4, pp. 377–430, 1933.

[28] Xiong Xiao, Xiaohai Tian, Steven Du, Haihua Xu, Eng Siong Chng, and Haizhou Li, "Spoofing speech detection using high dimensional magnitude and phase features: The ntu approach for ASVspoof 2015 challenge," in *Proc. Interspeech*, 2015.

[29] Xiaohai Tian, Zhizheng Wu, Xiong Xiao, Eng Siong Chng, and Haizhou Li, "Spoofing detection from a feature representation perspective," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2016.

[30] Leigh D Alsteris and Kuldip K Paliwal, "Short-time phase spectrum in speech processing: A review and some experimental results," *Digital Signal Processing*, vol. 17, no. 3, pp. 578–616, 2007.

[31] Michal Krawczyk and Timo Gerkmann, "STFT phase reconstruction in voiced speech for an improved single-channel speech enhancement," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 1931–1940, 2014.

[32] Bayya Yegnanarayana and Hema A Murthy, "Significance of group delay functions in spectrum estimation," *IEEE Transactions on Signal Processing*, vol. 40, no. 9, pp. 2281–2289, 1992.

[33] Steven M Kay, "The effects of noise on the autoregressive spectral estimator," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 27, no. 5, pp. 478–485, 1979.